UNITED STATES PATENT AND TRADEMARK OFFICE

UNITED STATES DEPARTMENT OF COMMERCE
United States Patent and Trademark Office
Address: COMMISSIONER FOR PATENTS
P.O. Box 1450
Alexandria, Virginia 22313-1450
www.uspto.gov

APPLICATION NO.	FILING DATE	FIRST NAMED INVENTOR	ATTORNEY DOCKET NO.	CONFIRMATION NO.
10/726,041	12/01/2003	Yasushi Saito	200300598-1	2665

22879 7590 05/08/2008

HEWLETT PACKARD COMPANY
P O BOX 272400, 3404 E. HARMONY ROAD
INTELLECTUAL PROPERTY ADMINISTRATION
FORT COLLINS, CO 80527-2400

EXAMINER	
STACY, BRENT S	

ART UNIT	PAPER NUMBER
2161	

NOTIFICATION DATE	DELIVERY MODE
05/08/2008	ELECTRONIC

Please find below and/or attached an Office communication concerning this application or proceeding.

The time period for reply, if any, is set in the attached communication.

Notice of the Office communication was sent electronically on above-indicated "Notification Date" to the following e-mail address(es):

JERRY.SHORMA@HP.COM
mkraft@hp.com
ipa.mail@hp.com

UNITED STATES PATENT AND TRADEMARK OFFICE

Commissioner for Patents
United States Patent and Trademark Office
P.O. Box 1450
Alexandria, VA 22313-1450
www.uspto.gov

BEFORE THE BOARD OF PATENT APPEALS AND INTERFERENCES

Application Number: 10/726,041
Filing Date: December 01, 2003
Appellant(s): SAITO ET AL.

Derek J. Westberg
For Appellant

EXAMINER'S ANSWER

This is in response to the appeal brief filed 3/7/2008 appealing from the Office action mailed 7/31/2007.

(1) Real Party in Interest

A statement identifying the real party in interest is contained in the brief.

(2) Related Appeals and Interferences

The examiner is not aware of any related appeals, interferences, or judicial proceedings which will directly affect or be directly affected by or have a bearing on the Board's decision in the pending appeal.

(3) Status of Claims

The statement of the status of claims contained in the brief is correct.

(4) Status of Amendments After Final

The appellant's statement of the status of amendments after final rejection contained in the brief is correct.

(5) Summary of Claimed Subject Matter

The summary of claimed subject matter contained in the brief is correct.

(6) Grounds of Rejection to be Reviewed on Appeal

The appellant's statement of the grounds of rejection to be reviewed on appeal is substantially correct. The changes are as follows:

Appellant's brief presents arguments relating to the drawings being properly objected. This issue relates to petitionable subject matter under 37 CFR 1.181 and not to appealable subject matter. See MPEP § 1002 and § 1201.

WITHDRAWN REJECTIONS

The following grounds of rejection are not presented for review on appeal because they have been withdrawn by the examiner. The previous 35 USC § 112 rejections regarding Claims 11 and 56-69 have been withdrawn.

(7) Claims Appendix

The copy of the appealed claims contained in the Appendix to the brief is correct.

(8) Evidence Relied Upon

20010044879	Moulton et al.	11-2001
2002/0107835	Coram et al.	8-2002

Popek, Gerald J. et al., "Replication in Ficus Distributed File Systems" Proceedings of the Workshop on Management of Replicated Data, (Nov. 1990), pp. 20-25

Zhang, Zheng et al., "Designing a Robust Namespace for Distributed File Services" Hewlett-Packard Laboratories (Oct. 2001), pp. 1-10

Yu, Haifeng et al. "The Costs and Limits of Availability for Replicated Services" Proceedings of the 18th ACM Symposium on Operating System Principles (SOSP) (Oct. 2001) pp. 1-14

(9) Grounds of Rejection

The following ground(s) of rejection are applicable to the appealed claims:

Claim Rejections - 35 USC § 112

In light of the Appellant's respective arguments or respective amendments, the previous 35 USC § 112 rejections to the claims have been withdrawn.

Claim Rejections - 35 USC § 103

The following is a quotation of 35 U.S.C. 103(a) which forms the basis for all obviousness rejections set forth in this Office action:

> (a) A patent may not be obtained though the invention is not identically disclosed or described as set forth in section 102 of this title, if the differences between the subject matter sought to be patented and the prior art are such that the subject matter as a whole would have been obvious at the time the invention was made to a person having ordinary skill in the art to which said subject matter pertains. Patentability shall not be negatived by the manner in which the invention was made.

Claims 1, 10, and 11 are rejected under 35 U.S.C. 103(a) as being unpatentable over "Replication in Ficus Distributed File Systems" (Popek et al.) in view of U.S. Patent Application Publication No. 2001/0044879 (Moulton, et al.).

For **Claim 1**, Popek teaches: "A method for a wide-area file system, including a plurality of nodes storing replicas of objects, the objects being files and directories, wherein for each replica of an object at a node, a parent directory for the object is replicated at the node, [Popek, page 22, 1st paragraph under "5 The Ficus Project"] the method comprising:

- propagating an update to a replica of the a file directory to other replicas of the file directory [Popek, page 21, 1st and 3rd paragraphs under "3 The Optimistic Model"]...and
- in response to receiving a propagated update to a replica of the file directory at a node, updating the replica for the file directory at the node" [Popek, page 21, 1st and 3rd paragraphs under "3 The Optimistic Model"].

Popek discloses the above limitations but does not expressly teach:

- "...via a graph, wherein each replica of the file directory has edges to only a subset of the other replicas such that all the replicas of the file directory are connected via the graph."

With respect to Claim 1, an analogous art, Moulton, teaches:

- "...via a graph, wherein each replica of the file directory has edges to only a subset of the other replicas such that all the replicas of the file directory are connected via the graph" [Moulton, Fig. 1].

It would have been obvious to one of ordinary skill in the art at the time of invention having the teachings of Moulton and Popek before him/her to combine Moulton with Popek because both inventions are directed towards replicating files across computers.

Moulton's invention would have been expected to successfully work well with Popek's invention because both inventions use computers on a network replicating files. Popek discloses Replication in Ficus Distributed File Systems (title) comprising updating and propagating replicas. However, Popek does not explicitly disclose a graph for propagating updates. Moulton discloses a system and method for distributed management of data storage (title) comprising a graph and way in which replicas are dispersed.

It would have been obvious to one of ordinary skill in the art at the time of invention having the teachings of Moulton and Popek before him/her to take the graph from Moulton and install it into the invention of Popek, thereby offering the obvious

advantage of having an efficient means of determining how the replicas replicate (Moulton, paragraphs [0032] and [0036]).

Claim 10 can be mapped to Popek (as modified by Moulton) as follows: "The method according to claim 1, wherein the replicas of the file directory include core replicas and non-core replicas, the parent directory for the file directory having edges only to the core replicas of the file directory and each core replica of the file directory having edges to the one or more of the non-core replicas of the file directory" [Moulton, paragraphs [0028] and [0078]].

Claim 11 can be mapped to Popek (as modified by Moulton) as follows: "The method according to claim 10, wherein in response to a user accessing an object at a node when no replica of the object exists at the node, the method further comprises steps of forming a non-core replica of the parent directory for the object at the node and forming a non-core replica of the object at the node" [Moulton, paragraphs [0078]].

Claims 2-9 and 56-66 are rejected under 35 U.S.C. 103(a) as being unpatentable over "Replication in Ficus Distributed File Systems" (Popek et al.) in view of U.S. Patent Application Publication No. 2001/0044879 (Moulton, et al.), further in view of "Designing a Robust Namespace for Distributed File Services" (Zhang et al.) (found in Appellant's IDS).

For **Claim 2**, Popek (as modified by Moulton) teaches: "The method according to claim 1, wherein."

Popek (as modified by Moulton) discloses the above limitation but does not expressly teach: "each replica of an object has a backpointer including an identification of a parent directory for the object and a name of the object in the parent directory.

With respect to Claim 2, an analogous art, Zhang, teaches: "each replica of an object has a backpointer including an identification of a parent directory for the object and a name of the object in the parent directory" [Zhang, page 2, 2nd paragraph above "2 Problem Abstraction" with Zhang, page 3, paragraph under "Back pointer"].

It would have been obvious to one of ordinary skill in the art at the time of invention having the teachings of Zhang and Popek (as modified by Moulton) before him/her to combine Zhang with Popek (as modified by Moulton) because both inventions are directed towards distributed file replication.

Zhang's invention would have been expected to successfully work well with Popek (as modified by Moulton)'s invention because both inventions use file replication. Popek (as modified by Moulton) discloses Replication in Ficus Distributed File Systems comprising updating and propagating replicas. However, Popek (as modified by Moulton) does not expressly disclose backpointers. Zhang discloses Designing a Robust Namespace for Distributed File Services comprising backpointers.

It would have been obvious to one of ordinary skill in the art at the time of invention having the teachings of Zhang and Popek (as modified by Moulton) before him/her to take the backpointers from Zhang and install it into the invention of Popek (as modified by Moulton), thereby offering the obvious advantage of minimizing overhead in guaranteeing namespace consistency and breaking down file service operations into

simple namespace primitives for easy namespace consistency and/or operation recovery.

Claim 3 can be mapped to Popek (as modified by Moulton and Zhang) as follows: "The method according to claim 2, wherein the parent directories are modified when the backpointer for a replica of an object at a node is not consistent with the parent directories for the replica of the object at the node" [Zhang, pages 2-3, "2 Problem Abstraction"].

Claim 4 can be mapped to Popek (as modified by Moulton and Zhang) as follows: "The method according to claim 3, wherein modifying the parent directories occurs only after a delay" [Zhang, pages 2-3, "2 Problem Abstraction"].

Claim 5 can be mapped to Popek (as modified by Moulton and Zhang) as follows: "The method according to claim 3, wherein multiple modifications to the parent directories at the node are performed according to an order in which corresponding updates occur" [Zhang, page 2, 2nd paragraph above "2 Prblem Abstraction" with Zhang, page 3 bullet 2 under "3 System Model and Failure Assumptions" with Zhang, page 4 paragraph under "4.2 Failure-free protocols" with Zhang, page 7 col. 1 with Zhang, page 7, Fig. 8].

Claim 6 can be mapped to Popek (as modified by Moulton and Zhang) as follows: "The method according to claim 3, wherein a modification is performed at the node and an earlier inconsistent modification is ignored" [Popek, page 21, middle paragraph in col. 2 or Zhang, page 3, paragraph above "3 System Model and Failure Assumptions"].

Claim 7 can be mapped to Popek (as modified by Moulton and Zhang) as follows: "The method according to claim 2, wherein a directory operation affects the backpointer for the object" [Zhang, page 4, paragraph under "Link"].

Claim 8 can be mapped to Popek (as modified by Moulton and Zhang) as follows: "The method according to claim 7, wherein the directory operation is selected from a group consisting of rename, link and unlink" [Zhang, page 2, section 2 "Problem Abstraction" with Zhang, page 4, Fig. 3 with Zhang, page 5, Fig. 4].

Claim 9 can be mapped to Popek (as modified by Moulton and Zhang) as follows: "The method according to claim 8, wherein when the backpointer for a replica of an object at a node is not consistent with the parent directories for the replica of the object at the node, further comprising modifying the parent directories to be consistent with the backpointer [Zhang, pages 2-3, "2 Problem Abstraction" with Zhang, pages 2-3, paragraphs under Table 2 to "3 System Model and Failure Assumptions"].

For **Claim 56**, Popek teaches: "A system including:

- a plurality of nodes that store replicas of objects, the objects being files and file directories, [Popek, page 22, 1st paragraph under "5 The Ficus Project"]
- wherein for each replica of an object at a node, the node stores a replica of a parent directory for the object [Popek, page 22, 1st paragraph under "5 The Ficus Project"] … and
- … and the nodes are configured to propagate updates to replicas of each file directory to other replicas of the file directory…" [Popek, page 21, 1st and 3rd paragraphs under "3 The Optimistic Model"].

Popek discloses the above limitations but does not expressly teach:

- "...and a backpointer having an identification of the parent directory for the object
- ...wherein each replica of a file directory has edges to only a subset of the other replicas of the file directory such that all the replicas of the file directory are connected via the graph
- ...via the graph."

With respect to Claim 56, an analogous art, Zhang, teaches:

- "...and a backpointer having an identification of the parent directory for the object" [Zhang, page 2, 2nd paragraph above "2 Problem Abstraction" with Zhang, page 3, paragraph under "Back pointer"].

With respect to Claim 56, an analogous art, Moulton, teaches:

- ...wherein each replica of a file directory has edges to only a subset of the other replicas of the file directory such that all the replicas of the file directory are connected via the graph [Moulton, Fig. 1]
- ...via the graph" [Moulton, Fig. 1].

It would have been obvious to one of ordinary skill in the art at the time of invention having the teachings of Zhang, Moulton and Popek before him/her to combine Zhang and Moulton with Popek because the inventions are directed towards distributed file replication.

Zhang's and Moulton's inventions would have been expected to successfully work well with Popek's invention because the inventions use file replication. Popek discloses Replication in Ficus Distributed File Systems comprising updating and

propagating replicas. However, Popek does not expressly disclose backpointers or explicitly disclose a graph for propagating updates. Zhang discloses Designing a Robust Namespace for Distributed File Services comprising backpointers. Moulton discloses a system and method for distributed management of data storage (title) comprising a graph and way in which replicas are dispersed.

It would have been obvious to one of ordinary skill in the art at the time of invention having the teachings of Zhang and Popek before him/her to take the backpointers from Zhang and to take the graph from Moulton and install them into the invention of Popek, thereby offering the obvious advantage of minimizing overhead in guaranteeing namespace consistency and breaking down file service operations into simple namespace primitives for easy namespace consistency and/or operation recovery. Using Moulton's graph offers the obvious advantage of having an efficient means of determining how the replicas replicate (Moulton, paragraphs [0032] and [0036]).

Claim 57 can be mapped to Popek (as modified by Moulton and Zhang) as follows: "The system according to claim 56, wherein in response to receiving a propagated update to a replica at a node, the node updates the parent directories for the file at the node" [Popek, page 21, paragraphs 3 and 4 under "3 The Optimistic Model" with Popek, page 22 paragraphs 1-3 under "5 The Ficus Project"].

Claims 58-66's limitation(s) have already been met by Claims 3-11's limitation(s), respectfully. Therefore, Claims 58-65 are rejected for the same reason(s) as stated above with respect to Claims 3-11, respectfully.

Claim 12 is rejected under 35 U.S.C. 103(a) as being unpatentable over "Replication in Ficus Distributed File Systems" (Popek et al.) in view of U.S. Patent Application Publication No. 2001/0044879 (Moulton, et al.), further in view of "The Costs and Limits of Availability for Replicated Services" (Yu et al.) (found in Appellant's IDS).

For **Claim 12**, Popek teaches: "The method according to claim 10, wherein."

Popek discloses the above limitation but does not expressly teach: "...a minimum number of core replicas are maintained according to a minimum replication factor."

With respect to Claim 12, an analogous art, Yu, teaches: "...a minimum number of core replicas are maintained according to a minimum replication factor" [Yu, page 1, last paragraph, with Yu, page 12, Figs. 11 and 12 with Yu, page 12, last paragraph in "5.3 Effects of Replication Scale" with Popek, page 21, 3rd paragraph under "3 The Optimistic Model"].

It would have been obvious to one of ordinary skill in the art at the time of invention having the teachings of Yu and Popek before him/her to combine Yu with Popek because both inventions are directed towards replicating files to achieve higher availability.

Yu's invention would have been expected to successfully work well with Popek's invention because both inventions use file replication. Popek discloses Replication in Ficus Distributed File Systems comprising updating and propagating replicas. However, Popek does not expressly disclose a minimum replication factor. Yu discloses the costs

and limits of availability for replicated services comprising determining a degree of replication.

It would have been obvious to one of ordinary skill in the art at the time of invention having the teachings of Yu and Popek before him/her to take the degree of replication from Yu and install it into the invention of Popek, thereby offering the obvious advantage of system developers being able to determine a degree of replication to achieve a target service availability.

Claims 13 and 15 are rejected under 35 U.S.C. 103(a) as being unpatentable over "Replication in Ficus Distributed File Systems" (Popek et al.) in view of U.S. Patent Application Publication No. 2001/0044879 (Moulton, et al.), further in view of U.S. Patent Application Publication No. 2002/0107835 (Coram et al.).

For **Claim 13**, Popek teaches: "The method according to claim 1."

Popek discloses the above limitation but does not expressly teach: "wherein a replica of an object is deleted by marking the replica as invalid."

With respect to Claim 13, an analogous art, Coram, teaches: "wherein a replica of an object is deleted by marking the replica as invalid" [Coram, paragraph [0047]].

It would have been obvious to one of ordinary skill in the art at the time of invention to combine Coram with Popek because both inventions are directed towards deleting data.

Coram's invention would have been expected to successfully work well with Popek's invention because both inventions use computers modifying computer storage.

Popek discloses Replication in Ficus Distributed File Systems comprising updating and deleting replicas, however Popek does not expressly disclose marking objects/replicas as invalid. Coram discloses a system and method for adaptive result set caching comprising marking objects/replicas as invalid.

It would have been obvious to one of ordinary skill in the art at the time of invention to take the marking objects/replicas as invalid from Coram and install it into the invention of Popek thereby offering the obvious advantage of having the objects/replicas still available after "deletion" (marking invalid) so that the objects/replicas may be recovered by a user if necessary.

Claim 15 can be mapped to Popek (as modified by Coram) as follows: "The method according to claim 13, further comprising periodically removing replicas marked as invalid" [Coram, paragraph [0050] with Coram, paragraph [0047]].

Claims 14 and 67-69 are rejected under 35 U.S.C. 103(a) as being unpatentable over "Replication in Ficus Distributed File Systems" (Popek et al.) in view of U.S. Patent Application Publication No. 2001/0044879 (Moulton, et al.), further in view of U.S. Patent Application Publication No. 2002/0107835 (Coram et al.), further in view of "Designing a Robust Namespace for Distributed File Services" (Zhang et al.) (found in Appellant's IDS).

For **Claim 14**, Popek (as modified by Coram) teaches: "The method according to claim 13."

Popek (as modified by Coram) discloses the above limitation but does not expressly teach: "wherein said marking the replica as invalid comprises removing the backpointer for the replica."

With respect to Claim 14, an analogous art, Zhang, teaches: "wherein said marking the replica as invalid comprises removing the backpointer for the replica" [Zhang, page 3, section 2 "Problem Abstraction"].

It would have been obvious to one of ordinary skill in the art at the time of invention to combine Zhang with Popek (as modified by Coram) because both inventions are directed towards distributed storage services with namespaces.

Zhang's invention would have been expected to successfully work well with Popek (as modified by Coram)'s invention because both inventions use files and directories with namespaces for distributed storage. Popek (as modified by Coram) discloses Replication in Ficus Distributed File Systems comprising updating and deleting replicas, however Popek (as modified by Coram) does not expressly disclose marking the replica as invalid comprises removing the backpointer for the replica. Zhang discloses the designing of a robust namespace for distributed file services comprising removing the backpointer for the replica.

It would have been obvious to one of ordinary skill in the art at the time of invention to take the removing the backpointer for the replica from Zhang and install it into the invention of Popek (as modified by Coram), thereby offering the obvious advantage of an efficient way of creating an orphaned object, thereby violating Zhang's

namespace rules and requiring special considerations for maintaining namespace integrity.

Claims 67-69's limitation(s) have already been met by Claims 13-15's limitation(s), respectfully. Therefore, Claims 67-69 are rejected for the same reason(s) as stated above with respect to Claims 13-15, respectfully.

(10) Response to Argument

In response to Appellant's argument that there is no suggestion to combine the references (Popek and Moulton), the examiner recognizes that obviousness can only be established by combining or modifying the teachings of the prior art to produce the claimed invention where there is some teaching, suggestion, or motivation to do so found either in the references themselves or in the knowledge generally available to one of ordinary skill in the art. See *In re Fine*, 837 F.2d 1071, 5 USPQ2d 1596 (Fed. Cir. 1988) and *In re Jones*, 958 F.2d 347, 21 USPQ2d 1941 (Fed. Cir. 1992). In this case, combining the prior arts offers the obvious advantage of having an efficient means of determining how the replicas replicate (Moulton, paragraphs [0032] and [0036]). I n the cited sections of Moulton, Moulton discusses Fig. 1. Fig. 1 of Moulton is a graph of computer nodes in a replication network of Moulton and how the nodes are all connected with each other. As combined, the graph shows the manner of which updates of Popek will go/travel to update other nodes. Without such a graph, it becomes inefficient to determine how replicas and updates to replicas will replicate (the use of a graph can also be seen in Popek, however, on p. 22, col. 2, ¾ down the col.).

As such, there is suggestion to combine the references found in the references themselves.

The Appellant's further argue that Popek and Moulton are not analogous arts because Popek allegedly teaches that "concepts taken from the field of databases and transaction-based data storage systems are not applicable to distributed file systems." Appellant's then at least imply that Moulton includes concepts taken from the field of databases and transaction-based data storage systems. However, Moulton is concerned with different ways of replicating files over a network and accessing the replicated, much like Popek. Moulton does not depend on a database to do this or any kind of transaction model/ transaction-based data storage system. As such Popek and Moulton appear to be readily combinable/modifiable for the rejections above.

Appellant's further argue that Moulton breaks a single data element into pieces and stores the pieces across the network nodes while Popek stores files across the network nodes. Appellant uses this argument as basis to try to show that the prior arts are allegedly further non-analogous. Even if Moulton does store pieces of files across network nodes, Moulton still teaches replicating the files across nodes (as the examiner stated in the final office action). This is because when a file is broken apart and replicated, the end result is the file being replicated. Additionally, RAIN level 1 in Moulton (see paragraph [0069]) supports pure mirroring of data (files) in Moulton. This is where files are replicated across nodes.

With respect to the Appellant's argument with respect to Claims 1 (and 56) for the prior art(s) allegedly not teaching "a plurality of nodes storing replicas of objects, the

object being files and file directories, wherein for each replica of an object at a node, a parent directory for the object is replicated at the node," the examiner respectfully disagrees. Popek, page 22, 1st paragraph under "5 The Ficus Project" was used to reject this limitation above. In the cited section, Popek teaches:

> "The system permits one to replicate files selectively within limits set by administrative control. That is, a collection of the file volume replicas are set up at various storage sites for a given logical file subtree. A given file may be replicated at any subset of the sites hosting a volume replica."

This citing teaches that files are replicated from "volume replicas." As such, these volumes hold files and at the very least can be considered folders (a.k.a. directories) containing files. Also, Popek, p. 21, paragraph 4 under "3 The Optimistic model" discussed the existence of a "directory system" in Ficus. These directories in Ficus are known to hold files (Popek, abstract, and knowledge generally known in the art). The volume alone can be considered a directory, the directories alone can be considered directories, or the directories in the volumes can be considered as a plurality of directories (with the volume being a parent directory of an object/file). These volume replicas with their files/file directories are the claimed objects being files and file directories. The replicas are set up at various storage sites. These storage sites are the claimed plurality of nodes storing replicas. The storage sites hosting the volume replicas means that the volume (folder/files) are copied/replicated at the node/storage site (for each storage site set up to host a volume replica). As such, a parent directory (volume) for the object (file within the volume) is replicated at the node (storage site). It

should be noted that Popek also explicitly teaches directories on Popek, p. 21, last paragraph. However, Popek teaches the claimed invention as claimed.

With respect to the Appellant's argument with respect to Claims 1 (and 56) for the prior art(s) allegedly not teaching "propagating an update to a replica of a file directory to other replicas of the file directory via a graph, wherein each replica of the file directory has edges to only a subset of other replicas such that all the replicas of the file direcotry are connected via the graph," the examiner respectfully disagrees. Popek, page 21, 1st and 3rd paragraphs under "3 The Optimistic Model" was used to teach the propagation of updates with Moulton, Fig. 1 mostly teaching the graph and the properties of the graph. Popek, page 21, 1st and 3rd paragraphs under "3 The Optimistic Model" teaches "...allow update to occur in at most one connected environment. Then propagate that update to other storage sites when communication is re-established." And "When multiple copies were re-connected, if any were out of date, the new version would be automatically propagated to make them current." Popek explicitly teaches propagating updates to other storage sites/nodes. These nodes, containing volume replicas (file directories) are seen as propagating updates to replicas of a volume/file directory to other replicas of the file directory. Popek does not go into detail explicitly how this is done, however, Moulton's graph can be applied to Popek's update propagation seeing that both inventions are dealing with networked computers, and therefore make a graph like Moulton, Fig. 1. Moulton, Fig. 1, shows a graph of networked nodes (details 105/106) connected by Local Area Networks (LANs, 104) and Wide Area Networks (WANs 103) and the Internet (102). The graph of Moulton is being

overlaid onto the invention of Popek such that the network in Popek looks like Moulton's network of Fig. 1. As such, the nodes in Moulton are Popek's storage sites. As can be seen in Moulton, Fig. 1, each replica of the file directory (node storing volume replica) has edges to only a subset of other replicas (e.g. edges only via LAN cloud 104 from node 105 to 105 or 106 within the LAN) such that all the replicas of the file directory (nodes storing volume replica) are connected via the graph (all nodes are eventually connected through the clouds in the graph). As such, the combination of reference teach the claimed limitations as claimed.

The Appellant did not appear to consider the combination of references used in the arguments in the appeal brief. As such the argument that "Popek does not discuss connections between replicas of file directories at all" is met, as discussed above, by Moulton. Appellant's argument that "Popek cannot disclose the particular manner of connecting file directories and the particular manner of propagating updates to the file directories" has been met with the explanation above. Appellant's argument that "there is no teaching or suggestion in Moulton that file directories might be connected in the particular manner recited" has also been met with the explanation above. Finally, the Appellant's argument that "there is no teaching or suggestion in Moulton...that updates to the file directories may be propagated" was met by Popek as can be seen above.

With respect to the Appellant's argument with respect to Claims 10 and 65 for the prior art(s) allegedly not teaching "wherein the replicas of the file directory include core replicas and non-core replicas, the parent directory for the file directory having edges to only to the core replicas of the file directory and each core replica of the file directory

having edges to one or more of the non-core replicas of the file directory," the examiner respectfully disagrees. Moulton, paragraphs [0028] and [0078] was used to best teach these claims. First, the examiner would like to discuss the interpretations of the claim. The examiner, in light of the specification (page. 26-27, lines 10-5), interprets core replicas as being replicas from which other replicas are made from, while non-core replicas are replicas made from core replicas. This interpretation appears to be a supported interpretation in the Appellant's specification (page. 26-27, lines 10-5). Moulton, paragraph [0078] teaches "In operation, every data write operation is executed to the primary node an all mirror nodes." This teaching shows that Moulton holds a node as being primary, while other nodes are mirror nodes. The primary node (inherently holding data like that of Popek (and can be seen in Moulton, paragraph [0078] reciting "primary image")) is the claimed "core replica" while the mirror nodes are the non-core replicas. The primary node nodes copies/replicas of objects across the network. As such, "replicas of the file directory include core replicas." The mirror nodes also hold replicas of the objects, and they are "replicas of the file directory include...non-core replicas." When an update occurs in Popek, the replica that originates the update (and propagates the update) is seen as a core replica since it is the replica from which other replicas are made. According to Fig. 1 of Moulton then, the update would propagate from one computer to another via the various networks (edges) available. As such, edges can be seen from a parent directory for the file directory having edges to only the core replicas of the file directory (the core update is propagated to non-core computers in the network via edges). Additionally, each core

replica of the file directory has edges to one or more of the non-core replicas of the file directory since the core update has to propagated to other replicas for file/directory consistency.

With respect to the Appellant's argument with respect to Claims 11 and 66 for the prior art(s) allegedly not teaching "wherein in response to a user accessing an object at a node when no replica of the object exists at the node, the method further comprises the steps of forming a non-core replica of the parent directory for the object at the node and forming a non-core replica of the object at the node," the examiner respectfully disagrees. Moulton, paragraph [0078] was used to reject this limitation. The cited section in Moulton teaches "Read operations attempt to first read the data from one of the nodes, and if that node is unavailable, a read from the mirror node is attempted." This teaches the reading method of Moulton (a user attempted to access an object at a node) that when a node is unavailable (when no replica exists) a non-core replica of the object and the object's parent directory will be created at the node. Moulton paragraph [0078] teaches about RAIN level 1 (similar to RAID level 1, but instead of being across storage devices it is across networked nodes). RAIN level 1 (and RAID level 1) guarantee mirroring of data between nodes (for RAIN) (while RAID is between storage devices). As such, when a replica is non-existent/not available, the system must recover from this failure in order to guarantee mirroring of data (Moulton [0078] teaches that RAIN level 1 "offers high reliability" and Moulton [0039] teaches that RAIN has fault tolerance and the ability to recover from such faults. As such, Moulton (with the prior modifications from Popek) teach the claimed limitations as claimed.

With respect to the Appellant's argument with respect to Claim 2 for the prior art(s) allegedly not teaching "wherein each replica of an object has a backpointer including an identification of a parent directory for the object and a name of the object in the parent directory," the examiner respectfully disagrees. Zhang, page 2, 2nd paragraph above "2 Problem Abstraction" with Zhang, page 3, paragraph under "Back pointer" was used to teach this limitation. In the cited sections, Zhang explictly teaches a "back pointer." Additionally, Zhang, page 3, paragraph under "Back pointer" teaches that back pointers are "references back to the parent directories of the object. In this cited paragraph, Zhang makes several references to Fig. 2(a) and Fig. 2(b) (located on p. 4). This figure show that the back pointer structure holds an identification of a parent directory for the object (site and inode # under "Parent dir reference"). Fig. 2(a) shows back pointers to parent directories for an object (dotted arrow lines). These parent directories are directories as defined by a "Directory entry structure" on p. 3in Zhang (Table 4). This directory entry structure includes an "object name." Therefore, a name of the object is in the parent directory. By using back pointers to the parent directory and parent directory entry structure the back pointer can also be seen as including an object name.

Even though Zhang appears to be directed toward distributed file services, Zhang has support (in a foot note on p. 2) for a replication service. Additionally, Zhang, p. 1, col. 2, lines 2-3 describe that DiFFS objects can be replicated (for a replication distributed file service), and Zhang, p. 1, col. 2, 2nd to last paragraph states that "this report is to investigate protocols for building robust namespaces in the context of

DiFFS. As such there is support for why one of ordinary skill in the art would look to Zhang to combine with Popek and Moulton. Also, as can be seen above, Zhang is at least reasonably pertinent to the particular problem with which the Appellant is concerned (consistency). Also, there is suggestion that Zhang is in the field of Appellant's endeavor and is reasonably pertinent to the particular problem with which the Appellant is concerned since an author of the Zhang reference is one of the inventors for the instant application (Christos Karamanolis). As such, Zhang is analogous prior art.

With respect to the Appellant's argument with respect to Claims 3, 9, 58, and 64 for the Zhang allegedly not suggesting "any actions that might be taken to detect whether there are inconsistencies among replicas of file or file directory, nor any actions that might be taken [in] response to detection of such an inconsistency," the examiner respectfully submits that this is not claimed subject matter in Claims 3, 9, 58, or 64. In pursuing compact prosecution, the examiner will not entertain arguments relating to not claimed subject matter. It should be noted that the Appellant has not argued subject matter of Claims 3, 9, 58, or 64 in the appeal brief. The Appellants have only submitted that Claims 3 and 58 are allegedly allowable at least because it depends from allegedly allowable claims 1 and 2 and 56, respectfully.

With respect to the Appellant's argument with respect to Claims 4 and 59 for the prior art(s) allegedly not teaching "a delay," the examiner respectfully disagrees. As can be seen throughout all the references, the references rely on networked computers. Sending data between computers always incurs a delay due to a numbers of factors

(e.g. processing power of the computers, network latency etc.). Even if the prior art(s) were implemented in real time and on one computer, electrical signals can only travel at (at most) the speed of light. This incurs delays of fractions of seconds. Either way, at least a delay is readily apparent since the references are implemented on computers in a computer network.

With respect to the Appellant's argument with respect to Claims 6 and 61 for the prior art(s) allegedly not teaching "wherein a modification is performed at the node and an earlier inconsistent modification is ignored," the examiner respectfully disagrees. Popek, page 21, middle paragraph in col. 2 or Zhang, page 3, paragraph above "3 System Model and Failure Assumptions" was used to rejection this limitation. Popek can be seen as teaching this limitation in that Popek, in the cited section teaches "When multiple copies were re-connected, if any were out of date, the new version would be automatically propagated to make them current." Here, Popek teaches that when a replication node re-connects, it will get the latest version of the file. This will occur no matter how many prior versions there were. These earlier, inconsistent modifications (updates) are ignored (since the node isn't connected). Alternatively, Zhang can also be seen as teaching this in Zhang, p. 3, paragraph above section 3. In this section Zhang teaches about inconsistencies where the object is not referenced and handling invalid references. Zhang, p. 7 col. 1 further expands on this in that "ii) the two generation#'s do not match. A NACK (non-acknowledge signal) is returned indicating that the back pointer has been already added by another link operation (LL(3) in Figure)." This indicates that inconsistent modifications are ignored. Further

inconsistent updates after this the Ficus system ignores the earlier inconsistent modification requests. As such, Popek and Zhang teach the claimed limitations as claimed.

In response to Appellant's argument that there is no suggestion to combine the references (Popek, Moulton, and Zhang) for Claim 56, the examiner recognizes that obviousness can only be established by combining or modifying the teachings of the prior art to produce the claimed invention where there is some teaching, suggestion, or motivation to do so found either in the references themselves or in the knowledge generally available to one of ordinary skill in the art. See *In re Fine*, 837 F.2d 1071, 5 USPQ2d 1596 (Fed. Cir. 1988) and *In re Jones*, 958 F.2d 347, 21 USPQ2d 1941 (Fed. Cir. 1992). In this case, combining the prior arts offers the obvious advantage of having an efficient means of determining how the replicas replicate (Moulton, paragraphs [0032] and [0036]). I n the cited sections of Moulton, Moulton discusses Fig. 1. Fig. 1 of Moulton is a graph of computer nodes in a replication network of Moulton and how the nodes are all connected with each other. As combined, the graph shows the manner of which updates of Popek will go/travel to update other nodes. Without such a graph, it becomes inefficient to determine how replicas and updates to replicas will replicate (the use of a graph can also be seen in Popek, however, on p. 22, col. 2, ¾ down the col.). Popek makes considerations for naming consistency in the Ficus system. Popek, p. 24, paragraph above "Acknowledgements" teaches:

> "A robust naming context system is also needed, which allows software and potentially replicated objects to be moved easily, without invalidating references (such as hypermedia links) in other objects. The name service must provide these transparency properties while at the

> same time being highly available, and hence replicated an mutually consistent."

Since Popek makes provisions for a naming context system, Zhang is combined with Popek to provide such naming context system. Zhang's naming system uses back pointers for the naming system, which teaches some limitation(s) of Claim 56. As such, there is suggestion to combine the references found in the references themselves.

The Appellant's further argue (for Claim 56) that Popek, Moulton, and Zhang are not analogous arts because Popek allegedly teaches that "concepts taken from the field of databases and transaction-based data storage systems are not applicable to distributed file systems." Appellant's then at least imply that Moulton includes concepts taken from the field of databases and transaction-based data storage systems. However, Moulton is concerned with different ways of replicating files over a network and accessing the replicated, much like Popek. Moulton does not depend on a database to do this or any kind of transaction model/ transaction-based data storage system. Zhang is similar to Popek in that it is providing a system for naming distributed/replicated data. Other similarities with Zhang and Popek were discussed above. As such Popek, Moulton, and Zhang appear to be readily combinable/modifiable for the rejections above.

Appellant's further argue (for Claim 56) that Moulton breaks a single data element into pieces and stores the pieces across the network nodes while Popek stores files across the network nodes. Appellant uses this argument as basis to try to show that the prior arts are allegedly further non-analogous. Even if Moulton does store

pieces of files across network nodes, Moulton still teaches replicating the files across nodes (as the examiner stated in the final office action). This is because when a file is broken apart and replicated, the end result is the file being replicated. Additionally, RAIN level 1 in Moulton (see paragraph [0069]) supports pure mirroring of data (files) in Moulton. This is where files are replicated across nodes.

With respect to the Appellant's argument with respect to Claim 56 for the prior art(s) allegedly not teaching "wherein for each replica of an object at a node, the node stores a replica of a parent directory for the object and a backpointer having an identification of the parent directory for the object," the examiner respectfully disagrees. As shown and combined above, "wherein for each replica of an object at a node, the node stores a replica of a parent directory for the object" is taught by Popek, page 22, 1st paragraph under "5 The Ficus Project" while "and a backpointer having an identification of the parent directory for the object" is taught by Zhang, page 2, 2nd paragraph above "2 Problem Abstraction" with Zhang, page 3, paragraph under "Back pointer." It appears the Appellant has not considered the combination of the references teaching the limitations of the claims. "wherein for each replica of an object at a node, the node stores a replica of a parent directory for the object" was shown to be taught above with connection to the response to arguments regarding Claim 1. "and a backpointer having an identification of the parent directory for the object" was shown to be taught above with connection to the response to arguments regarding Claim 2. See above.

With respect to the Appellant's argument with respect to Claim 57 for the prior art(s) allegedly not teaching "wherein in response to receiving a propagated update to a replica for the file directory at a node, the node updates the parent directories for the file at the node," the examiner respectfully disagrees. Popek, page 21, paragraphs 3 and 4 under "3 The Optimistic Model" with Popek, page 22 paragraphs 1-3 under "5 The Ficus Project" was used to reject Claim 57. Popek, p. 21, cited paragraph 3 teaches "When multiple copies were re-connected, if any were out of date, the new version would be propagated to make them current." Popek explicitly teaches propagating updates (files or folders). The parent directories leading up to the update are updated (contents of the directory/parent directories are updated so that users can access correct appropriate data). Popek, p. 21, cited paragraph 4 teaches that "one might argue that the only updates done to directories are record inserts and deletes" while Popek, p. 22, cited paragraph 3 "guarantee[s] all updates and deletions to a partially replicated directory structure." As such, updates are done to parent directories for the file at the node when a propagated update to a replica for a file directory is received.

With respect to the Appellant's argument with respect to Claim 12 for the prior art(s) allegedly not teaching "wherein a minimum number of core replicas are maintained according to a minimum replication factor," the examiner respectfully disagrees. As discussed above, "core replicas" were taught above in connection with the examiner's response to the arguments regarding Claims 10 and 65. Yu, page 1, last paragraph, with Yu, page 12, Figs. 11 and 12 with Yu, page 12, last paragraph in "5.3 Effects of Replication Scale" with Popek, page 21, 3rd paragraph under "3 The

Optimistic Model" was used to reject this claim. The cited sections in Yu are teaching how to determine a best "degree of replication." Maintaining a degree of replication is maintaining a minimum number of replicas. Any degree of replication is replication. Also, p. 12 of Yu teaches that "the optimal number of replicas tops out at 2." This teaches that at least one, bur preferably 2 replicas are preferred. Yu, p. 12 shows figures where the x axis is a number of replicas (degree of replication as described in the paragraph describing Fig. 11). 1 on the x axis is 1 replica. With 1 replica there is 1 original file (for example) and 1 replica. 2 replicas is 2 replicas and 1 original file. This is showing that in Yu, at least 1 (minimum) original (core) replica is being maintained to make the replica files. Popek was also cited in rejecting this claim since the citing in Popek at least suggests providing at least 1 (minimum replication factor) replica in teaching "Imagine an architecture in which update were allowed whenever a copy of the needed data were available" (emphasis added). As such, Yu, and at least Popek teach "wherein a minimum number of core replicas are maintained according to a minimum replication factor."

With respect to the Appellant's argument with respect to Claims 13 and 67 for the prior art(s) allegedly not teaching "wherein a replica of an object is deleted by marking the replica as invalid," the examiner respectfully disagrees. Coram teaches a result set cache that holds data from a source (e.g. database). The data in the cache is a replica of the data in the database until the data is invalid (doesn't match the database data). The cache in Coram speeds up the system in Coram since accessing a database can be a timely operation (Coram, paragraph [0008]). Coram, paragraph [0047] was used

to rejection Claims 13 and 67. Coram, paragraph [0047] teaches that result sets in the cache become invalidated when they are no longer considered valid. These "Invalid result set can be removed from memory immediately or simply replaced in cache memory by one or more valid result sets as they become available." As such, Coram teaches "wherein a replica of an object is deleted by marking the replica as invalid."

Appellant's further argue that Coram is not analogous prior art and there is no suggestion/motivation to combine Coram with Popek since Popek allegedly teaches "that concepts taken from the filed of database and transaction-based data storage systems are not applicable to distributed file systems."

In response to Appellant's argument that Coram is nonanalogous art, it has been held that a prior art reference must either be in the field of Appellant's endeavor or, if not, then be reasonably pertinent to the particular problem with which the Appellant was concerned, in order to be relied upon as a basis for rejection of the claimed invention. See *In re Oetiker*, 977 F.2d 1443, 24 USPQ2d 1443 (Fed. Cir. 1992). In this case, Coram indirectly teaches replicas of data objects (Coram is in the field of Appellant's endeavor). Also, Coram is reasonably pertinent to the particular problem with which the Appellant was concerned since Coram teaches invalidating replicated data and deleting it (see above).

In response to Appellant's argument that there is no suggestion to combine the references, the examiner recognizes that obviousness can only be established by combining or modifying the teachings of the prior art to produce the claimed invention where there is some teaching, suggestion, or motivation to do so found either in the

references themselves or in the knowledge generally available to one of ordinary skill in the art. See *In re Fine*, 837 F.2d 1071, 5 USPQ2d 1596 (Fed. Cir. 1988) and *In re Jones*, 958 F.2d 347, 21 USPQ2d 1941 (Fed. Cir. 1992). In this case, suggestion to combine the references is found with the knowledge generally available to one of ordinary skill in the art. Appellant appears to refer to Popek p. 23 2nd paragraph under section 6 as Popek teaching away from combining the references as above. Popek explicitly teaches "We argue that one should not even apply transaction models to collection of copies of an individual object, as that means update (and even, in some cases, reading) is impossible if any particular copy is missing." This teaching is talking about how transaction models should not be used to aid in consistency, synchronization and naming (see Popek, p. 23, 1st paragraph under section 6). Coram is not concered with a transaction model for synchronization, concurrency, and/or naming. As discussed above, Coram is concerned with result set caching to improve access time to data. Speeding up the invention of Popek is a reasonable suggestion for one of ordinary skill in the art at the time of invention to combine the references.

With respect to the Appellant's argument with respect to Claims 14 and 68 for the prior art(s) allegedly not teaching "wherein said marking the replica as invalid comprises removing the backpointer for the replica," the examiner respectfully disagrees. Zhang, page 3, section 2 "Problem Abstraction" was used to reject this limitation. This section teaches about the link and unlink operations that are used "to every namespace operation." The unlink protocol is further described on Zhang, p. 5, and Zhang, Fig. 4. Zhang on p. 5 teaches "the reference is first removed from the namespace (directory),

before this fact is reflected on the referenced object (back-pointer removal)." This teaches that when an object (like a directory) is removed, the back pointer is also removed. Since Zhang along with its combined art(s) are dealing with replicated files, the backpointer is removed for the replica.

The other claims argued merely because of a dependency on a previously argued claim(s) in the brief presented to the examiner, filed March 7th, 2008, are moot in view of the examiner's interpretation of the claims and art and are still considered rejected based on their respective rejections from at least a prior Office action (part(s) of recited again above).

(11) Related Proceeding(s) Appendix

No decision rendered by a court or the Board is identified by the examiner in the Related Appeals and Interferences section of this examiner's answer.

For the above reasons, it is believed that the rejections should be sustained.

Respectfully submitted,

/BRENT STACE/

Examiner, Art Unit 2161

Conferees:

/Apu M Mofiz/

Supervisory Patent Examiner, Art Unit 2161

/Eddie C Lee/

Supervisory Patent Examiner, TC 2100